arXiv:1604.04764v1 [cs.NE] 16 Apr 2016

Closed loop interactions between spiking neural network and robotic simulators based on MUSIC and ROS

Philipp Weidel[1*], Mikael Djurfeldt[2,1], Renato Duarte[1,3,4], Abigail Morrison[1,4,5,6]

[1]Institute for Advanced Simulation (IAS-6),
Theoretical Neuroscience &
Institute of Neuroscience and Medicine (INM-6),
Computational and Systems Neuroscience &
JARA BRAIN Institute I
Jülich Research Center and JARA
Jülich, Germany

[2]PDC Center for High Performance Computing
KTH
Stockholm, Sweden

[3]Faculty of Biology
Albert-Ludwig University of Freiburg
Freiburg im Breisgau, Germany

[4]Bernstein Center Freiburg
Albert-Ludwig University of Freiburg
Freiburg im Breisgau, Germany

[5]Simulation Laboratory Neuroscience –
Bernstein Facility for Simulation and Database Technology
Institute for Advanced Simulation
Jülich Aachen Research Alliance
Jülich Research Center
Jülich, Germany

[6]Institute of Cognitive Neuroscience
Faculty of Psychology
Ruhr-University Bochum
Bochum, Germany

*Forschungszentrum Jülich GmbH
Institute of Neuroscience and Medicine (INM-6)
52425 Jülich, Germany
e-mail: p.weidel@fz-juelich.de

Abstract

In order to properly assess the function and computational properties of simulated neural systems, it is necessary to account for the nature of the stimuli that drive the system. However, providing stimuli that are rich and yet both reproducible and amenable to experimental manipulations is technically challenging, and even more so if a closed-loop scenario is required. In this work, we present a novel approach to solve this problem, connecting robotics and neural network simulators. We implement a middleware solution that bridges the Robotic Operating System (ROS) to the Multi-Simulator Coordinator (MUSIC). This enables any robotic and neural simulators that implement the corresponding interfaces to be efficiently coupled, allowing real-time performance for a wide range of configurations. This work extends the toolset available for researchers in both neurorobotics and computational neuroscience, and creates the opportunity to perform closed-loop experiments of arbitrary complexity to address questions in multiple areas, including embodiment, agency, and reinforcement learning.

1 Introduction

Studying a functional, biologically plausible neural network that performs a particular task is highly relevant for progress in both neuroscience and robotics. The major focus on this topic in the field of robotics consists of using of neural networks of varying degrees of complexity for controlling robots. So far, the majority of research has focused on non-spiking, artificial neural networks for this task (Quiñonez et al., 2015; Antonelo et al., 2007; Dasgupta et al., 2013), but there is considerable interest in investigating the capacities of spiking neural networks.

In computational neuroscience, the study of simulated neural networks is paramount to gain a better understanding of the processes underlying learning/adaptation to complex environments and global behavior (e.g. Chorley & Seth, 2008). However, neural network simulators do not typically include functionality for representing environments and sensory input. As a consequence, most tasks used to test the function of a simulated neural network are hardcoded to represent a highly specific task. This has the disadvantage that virtual experiments are complex and time consuming to develop and adapt. More importantly, tasks defined in this way are rather artificial (Jitsev et al., 2012; Frémaux et al., 2013; Potjans et al., 2011; Legenstein & Maass, 2014; Friedrich & Lengyel, 2015). Whereas there is certainly value in investigating very simplified tasks and sensory representations, it is also vital to be able to check that proposed neural architectures are capable of handling richer, noisier and more complex scenarios.

However, design of environments, representation of sensory states and conversion of motor commands into movements are primary features of robotic simulators. Therefore, it would be of great value to both research fields if the powerful tools developed for robotic simulation and spiking neural network simulation could be made to work together. This would allow researchers from both fields to perform simulated closed loop experiments with flexible experiment design, rich sensory input, complex neuronal processing, and motor output. The major challenge lies in the fact that robotic simulators communicate via continuous data streams, while neural network simulators communicate with spike events (so as their biological counterparts). Thus, a principled approach for bi-directional conversion is required.

So far there have been few attempts to address this. Moren et al. (2015) describe a technical setup for a specific use case, in which a robot is connected to a neural network model of early saccade movement simulated with the Neural Simulation Tool (NEST: Gewaltig & Diesmann, 2007). As the robot is not in the same location as the cluster where the neural simulation is running, the connection is realized via the internet using an SSH tunnel and the Multi-Simulation Coordinator (MUSIC: Djurfeldt et al., 2010), a middleware facilitating communication between different neural simulators in one simulation. Using a minimal model and optimal number of cores, the NEST simulation ran a factor of two slower than real time and achieved an output response of $110 - 140\,\mathrm{ms}$. Although this study provides a proof of concept for interaction between robotics and neural simulators, it does not represent a general solution, in part due to the limited

implementation detail provided.

A more general solution was presented by Hinkel et al. (2015), who describe an interface between the Robotic Operating System (ROS: Quigley et al., 2009) and NEST. ROS is the most popular middleware for the robotic community, with interfaces to many robotic simulators and also robotic hardware. It allows the user to create their own simulated robots in great detail using the Unified Robot Description Format. Similarly, NEST is one of the most commonly used neural simulators for spiking point neuron models. However, the proposed interface is on the basis of Python: the neural network must be simulated in brief periods in an external loop, reading out and communicating the spikes to the robotic simulator at each iteration. Although no performance data are provided, the overheads inherent in repeatedly stopping and starting NEST imply severe performance limitations for this approach.

In this paper we present an alternative approach based on ROS and MUSIC that is both general and efficient, enabling real-time performance for up to hundreds of thousands of neurons on our hardware. Our approach provides roboticists the opportunity to extend their work on neural control into the realm of spiking neural networks using any neural simulator implementing a MUSIC interface, including NEST or NEURON (Hines & Carnevale, 2001). Conversely, our toolchain frees the researcher from the constrictions of the hardcoding approach described above, by enabling neural simulators to be connected to any of the myriad robotic simulators implementing a ROS interface (including Gazebo[1], Morse[2] or Webots[3]). To demonstrate the capabilities of the toolchain, we implement a Braitenberg Vehicle (Braitenberg, 1986) in which the agent is simulated in Gazebo and the neural controller in NEST.

1.1 Description of the toolchain

The main purpose of the present toolchain is to capture a broad set of use cases by connecting robotic with neural simulators in a generic way, obtained by interfacing the two middlewares (ROS and MUSIC). The interface is licensed under GPLv3 and available on Github[4]. In addition to our own repository, we plan to distribute the toolchain as a plug-in for MUSIC[5].

In this work, we propose a set of possible solutions to overcome the problem of converting between continuous and spiking signals. In particular, we investigate the performance of three different kinds of encoders (NEF, regular rate coding, Poisson rate coding) and a linear readout decoder. As the correct encoding mechanism is debatable and dependent on the scientific question being addressed, in addition to our own mechanisms, we provide extensibility such that researchers can implement their own custom encoders, decoders and adapters in Python or C++.

[1] http://gazebosim.org/
[2] https://www.openrobots.org/wiki/morse/
[3] https://www.cyberbotics.com
[4] https://github.com/weidel-p/ros_music_adapter
[5] https://github.com/INCF/MUSIC

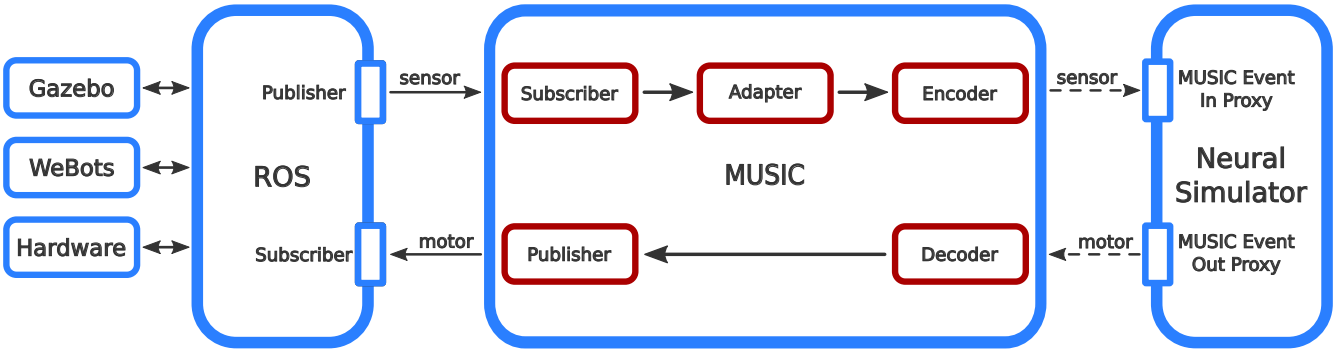

Figure 1: Information flow in the ROS-MUSIC toolchain. The boxes with red outlines depict our extensions to MUSIC that facilitate the bidirectional conversion of continuous and spiking signals (solid and dashed arrows, respectively)

To achieve these ends, our interface to ROS extends MUSIC by three different kind of binaries: adapters, encoders and decoders (see Figure 1). Encoders and decoders are used to translate between spiking data and continuous data while adapters can be used for pre-processing data and connect to ROS. All binaries run in their own process and solve only one specific purpose. This way, the interface is highly modular and the implementation of custom adapters, encoders or decoders is as simple as possible.

2 Materials and Methods

2.1 The MUSIC configuration file

The adapters, encoders, decoders and simulators introduced in the toolchain are specified, connected and parameterized in the MUSIC configuration file. Detailed information about the MUSIC configuration file can be found in the MUSIC manual[6] and examples of complete configuration files for this toolchain can be found in the Github repository.

2.2 Encoding, decoding and adaptation

The choice of encoder and decoder to convert continuous to spiking signals is non-trivial, and potentially dependent on the scientific question or task. Here, we describe three possible encoding mechanisms and one decoding mechanism. We also describe a simple adapter to map the dimenstionality of the input signal to the dimensionality of the receiving network.

[6] http://software.incf.org/software/music

2.2.1 Regular rate coding

The simplest way to encode continuous data into spiking data is regular rate coding. i. e. mapping the input signal to a time-dependent firing rate variable. Each input dimension is encoded by a single neuron by transforming the input signal, which is a continuous signal in the range $[-1, 1]$, into a time-dependent variable representing the spike density of the encoding neuron. This is achieved by calculating the interspike interval of each encoding neuron at time t by:

$$ISI_n(t) = \frac{1}{v_{\min} + (v_{\max} - v_{\min})\frac{1+I_n(t)}{2}} \tag{1}$$

where $I_n(t)$ is the input which is mapped to neuron n. The firing rate is scaled between $[v_{\min}, v_{\max}]$ which are free parameters of the regular rate encoder. This encoding mechanism leads to very regular spike patterns but has the clear advantage of being computationally very efficient. Note that this approach does not take into consideration biological neuron properties such as refractoriness, but these could be included without much impact on the computational efficacy.

2.2.2 Poisson rate coding

In this encoding scheme, we follow a similar implementation to that described in the previous section, but introduce a stochastic component in the generated spike trains, by transforming the input signal into a time-dependent variable representing the intensity of a stochastic point process with Poissonian statistics. This can be achieved by using the inverse of the expression (Equation 1) for interspike interval in the previous section as the rate parameter for an exponential ISI distribution:

$$ISI_n(t) = -\frac{1}{v_{\min} + (v_{\max} - v_{\min})\frac{1+I_n(t)}{2}} \ln(r) \tag{2}$$

where r is a random number drawn from a uniform distribution in $[0, 1)$ and $I_n(t)$ is the input mapped to neuron n. The firing rate is scaled between $[v_{\min}, v_{\max}]$ which are free parameters of the Poisson rate-encoder. This encoding mechanism produces spike trains with Poissonian statistics, similar to those typically assumed to occur in central neurons in the mammalian brain (see, for example, Calvin & Stevens (1967); Dayan & Abbott (2001) or Averbeck (2009) for a counter-argument), and therefore are thought to be a more realistic approximation than regular spike trains. As Section 3 shows, the computational complexity of this approach is comparable to that of regular rate encoding.

2.2.3 Neural Engineering Framework

The Neural Engineering Framework (NEF: Eliasmith & Anderson, 2004) is a generic formalism to represent stimuli with neural ensembles that allows a wide variety of functions to be realised by deriving optimal projections between neural populations. In the simplest formulation, NEF encodes a vector of continuous data I as an input current to an integrate-and-fire neuron whose activity is defined by

$$a_n(I) = G_n[\alpha_n \cdot I + I_n^{\text{bias}}] \tag{3}$$

where G represents a generic non-linear functional which is determined by the neuron dynamics, α_n is the neuron's response preference, or tuning curve, and I_n^{bias} a constant bias current setting the base activation of neuron n. It is worth noting that this encoding does not depend on the specific form of G and any neuron model can be used for this operation. The original signal I can be adequately reconstructed (depending on the properties of G and alpha) by linearly combining the activity a of the encoded representations

$$\hat{I} = \sum_n a_n(I)\phi_n \tag{4}$$

where the decoding weights ϕ are obtained by linear regression in order to minimize the reconstruction error.

2.2.4 Linear decoder

Having described possible options explored in this work to convert continuous signals into discrete spike event trains, we now explore the reverse operation. To transform the spiking activity of a given neural ensemble to a continuous signal, which can be used, for example, to provide motor commands to the simulated robot, we first perform a low-pass filter of each spike train by convolving with a causal (i.e. only defined for $t \geq 0$) exponential kernel $k(t) = \exp(-t/\tau_{\text{dec}})$ with time constant τ_{dec}

$$a_n(t) = \sum_i k(t - t_{i,n}) \tag{5}$$

where $t_{i,n}$ represents the i-th spike time of neuron n. The resulting activities at time t can be linearly combined to obtain a continuous output signal

$$z_k = \sum_n a_n(t)\phi_{nk} \tag{6}$$

where weights ϕ_{nk} define the contribution of ensemble neuron n to readout unit k.

2.2.5 Signal adapter

For the experiments carried out below we implement a simple adapter that maps the m dimensions of the input signal to the n dimensions of the receiving neurons. Note that this is only necessary for regular and Poisson rate encoding; the NEF encoder already incorporates this functionality by construction. For the performance measurements, each receiving neuron receives all dimensions of the input signal. For the Braitenberg vehicle, the input signal is split up into two hemispheres, each mapped to one of the two controlling neurons. For more complex experiments, a more sophisticated adapter can be implemented.

2.3 Performance Measurements

In order to be as close as possible to a real use-case, we simulate the whole toolchain while measuring the performance capabilities of the different parts. The agent used for the performance measurements is a four wheeled, mobile robot simulated in Gazebo with an attached virtual laser scanner. The laser scanner has 100 beams with an update rate of $20\,\mathrm{Hz}$ and a maximal range of $5\,\mathrm{m}$. We use the `SkidSteerDrivePlugin` provided by Gazebo, which allows us to steer the robot with a `ROS::Twist` message, updated with a rate of $20\,\mathrm{Hz}$. During the measurements involving the NEF encoder, we keep the parameters of the integrate-and-fire (IAF) neurons unchanged and simulate these neurons with a resolution of $1\,\mathrm{ms}$ using exact integration (Rotter & Diesmann, 1999; Morrison et al., 2007; Hanuschkin et al., 2010). For decoding, we always use a linear readout where the spiking activity is filtered with an exponential kernel with a resolution of $1\,\mathrm{ms}$. The MUSIC adapters run with an update rate equal to the update rate of the sensory input and motor command output. The measurements were executed on one node of a transtec CALLEO 551 cluster and were averaged over five trials with a simulation-time of ten seconds. Measurements involving neural simulators were carried out with NEST 2.8 (Eppler et al., 2015) and NEURON v6.2.

2.3.1 Real-time factor

The sensory input, the motor output and the processing of the sensory data are all updated in parallel but asynchronously. For robotic applications it is crucial that the execution time for processing input data in each time step is less than or equal to the wall-clock time of each time step, meaning the process is running in real-time.

The real-time factor (RTF) is calculated by dividing the simulated time, t_{sim}, by the wallclock time required to run the simulation, t_{run}.

$$RTF = \frac{t_{\text{sim}}}{t_{\text{run}}} \tag{7}$$

The total time the toolchain needs to perform a simulation can be divided into different components

$$t_{\text{total}} = t_{\text{build}} + t_{\text{run}} + \epsilon \tag{8}$$

where t_{build} is the time the toolchain needs for initialization (allocate memory etc), t_{run} is the actual time the toolchain needs to perform the simulation and ϵ is the time which is not covered by t_{build} and t_{run} (i.e. garbage collection, freeing memory, etc.). In our measurements we clearly separate the buildup phase from the run-time phase by synchronizing the processes after the initialization of the software using an MPI barrier. This allows us to measure t_{run} consistently over all processes.

We use this measurement to investigate different aspects of the toolchain's performance, as described below.

Dimensionality of the input The sensory input to a robot can be very complex and high dimensional. The number of neurons needed to encode a multidimensional input strongly depends on the encoding mechanism. In NEF, about 100 IAF neurons per dimension are used when encoding a stimulus in order to be able to decode that stimulus with a root-mean-square error about 1% (Eliasmith & Anderson, 2004). Using rate coding, the number of neurons can be freely chosen. In the example of the Braitenberg Vehicle (Braitenberg, 1986), we use only two neurons for encoding the complete sensory input (see Section 3.1).

To examine the scalability of the sensory input in our toolchain, we measure how many encoding neurons for different encoders we can simulate until the real-time performance breaks down. We use a binary search to find these limits in an efficient way. When measuring the performance with NEST and NEURON, each neuron instantiated by the encoder is matched by a corresponding neuron in the neuronal simulator, which receives the input spike train and repeats it to the decoder as an identical spike train, thus implementing a minimal processing network. In NEST this is realized by the `parrot_neuron`, a neuron model that emits a spike for each spike received. In NEURON the repeater neurons are of type `IntFire1` equipped with very high input synaptic weights and zero refractory time, thus enforcing the neuron to spike after each input spike. For the rate and Poisson encoder, the encoder neurons have a low firing rate between 1 and $2\,\text{Hz}$. Except for NEST (7 processes) and NEURON (40 processes), each process in the toolchain runs on one dedicated process.

Bandwidth Apart from the computational limitations of simulating neurons, communication is another potential bottleneck in this toolchain. Parts of the toolchain are communicating action potentials (events) between different processes (i.e. from encoder to NEST). In order to investigate the influence of the communication on the real-time performance, we use a rate based encoder and measure the real-time factor as a function of the firing rate of the encoding neurons. We choose the amount of neurons close to the border of real-time capability, so that we can see an immediate impact of the firing-rate on the performance of the toolchain.

Latency The latency between sensory input and motor output can be crucial for the robotic application. A fast reaction time is needed for many applications (for example, catching a ball). In order to measure the latency (or reaction time), we artificially disturb the sensory signal of the robotic simulator, switching discontinuously from the minimum to the maximum sensor range. A very simple encoder responds to this change by beginning to produce spikes, which the neural simulator receives over its MUSIC Event In Proxy and repeats its the MUSIC Event Out Proxy (see Figure 1). A decoder responds to the arrival of this spike train by producing a motor command, which is conveyed over ROS to the robotic simulator. The latency of the tool chain is therefore the wall-clock time between the change in the sensory signal and the reception of the motor signal.

Overhead For measuring the input scalability as described above, we use a MUSIC time step equal to the ROS update rate. This minimizes the communication overhead as only new data is communicated. As latency depends on the MUSIC time step, we additionally investigate the communication overhead of the toolchain. To do this, we determine the real-time factor of the toolchain, using a simple rate encoding mechanism, whilst systematically varying the MUSIC time step and the number of simulated neurons.

3 Results

Figure 2 shows the effect of the computational load of the different encoding mechanisms on the real-time factor, with and without neural simulators involved in the toolchain. For the rate or Poissonian encoding mechanisms (Figure 2**a,b**), simulations of up to $150,000$ encoding neurons are possible in real-time when using NEST or when the toolchain does not include a neural simulator. When using NEURON, the real-time factor breaks down at about 1000 neurons. For the NEF encoding mechanism (Figure 2**c**), simulations of up to $20,000$ encoding neurons are possible in real-time when using NEST or no neural simulator in the toolchain. In this case too, the performance when using NEURON breaks down at about 1000 neurons.

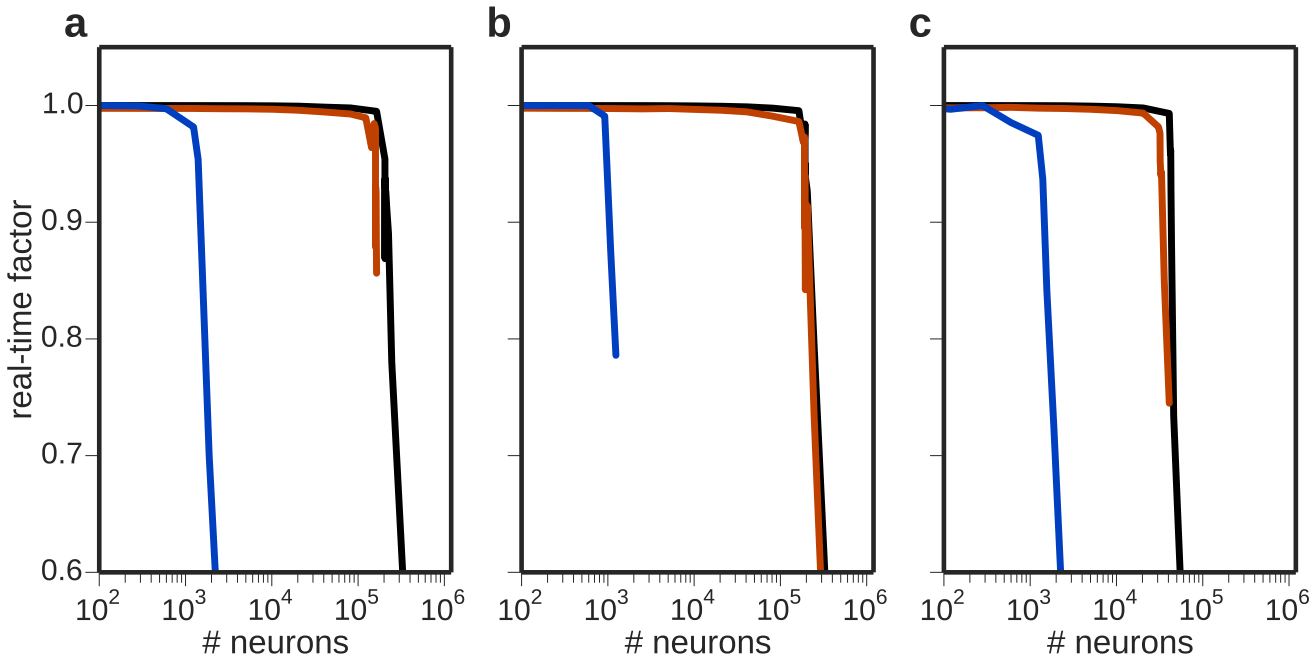

Figure 2: Scalability of the toolchain. Real-time factor as a function of the number of neurons used for encoding using different encoders: **a)** rate encoding mechanism; **b)** Poissonian encoding mechanism; **c)** NEF encoding mechanism. In each panel, the black curve shows the real-time performance in the absence of a neuronal simulator, the red curve shows the performance using NEST and the blue curve the performance using NEURON.

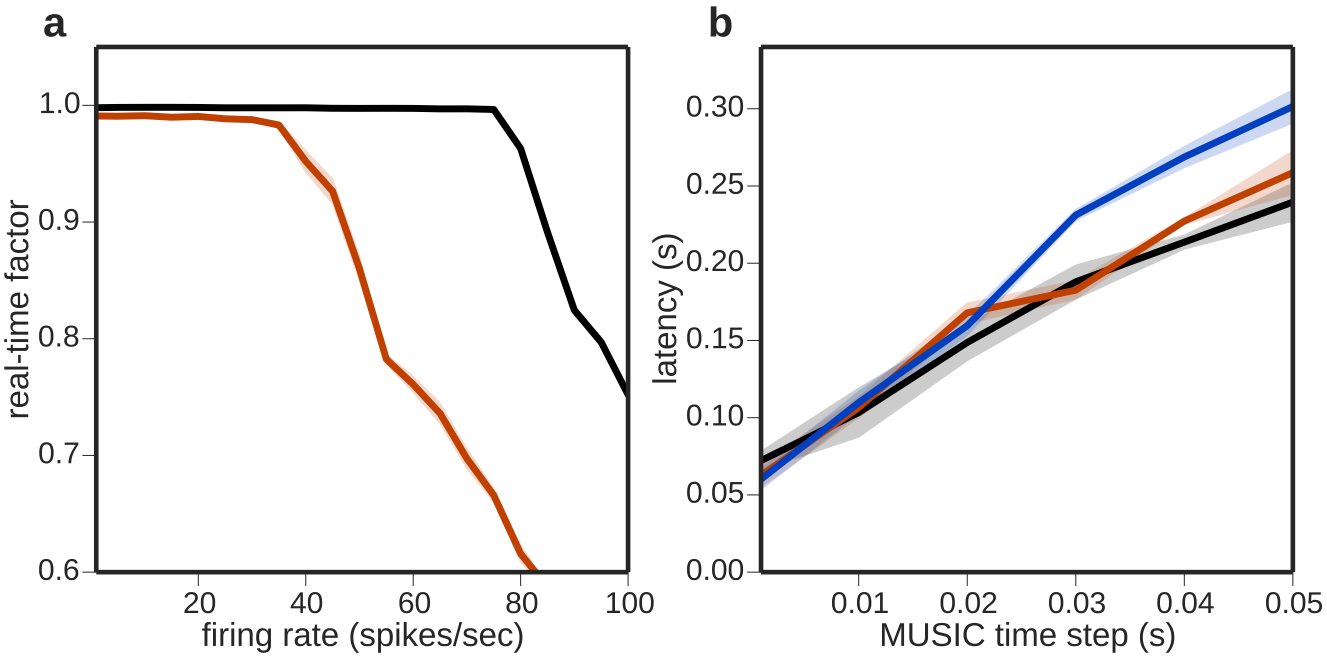

Figure 3: Bandwidth limitations of the toolchain. **a)** Real-time factor as a function of neuronal firing rate per neuron for a regular rate encoder of $50,000$ neurons without a neural simulator (black curve) and with NEST (red curve). **b)** Latency as a function of MUSIC time step without a neural simulator (black curve), with NEST (red curve) and with NEURON (blue curve).

These results demonstrate that the NEF encoder is computationally more expensive than either the rate or Poissonian encoders, and that NEURON is computationally expensive in this context, limiting the real-time performance to 1000 neurons regardless of which encoding mechanism is used.

The breakdown of the real-time performance of the regular rate and Poissonian encoders happens at about the same amount of neurons, which leads to the question what the actual bottleneck of the toolchain is in these cases. Computational complexity is only one limiting factor in this toolchain; another candidate is the communication.

But as Figure 3**a** shows, the real-time factor breaks down at about $40\,\mathrm{Hz}$ using $50,000$ encoding neurons. In other words, the toolchain is able to communicate about $2,000,000$ spikes per second, which is more than was required for the simulations measured in Figure 2. Thus we can conclude firstly that communication was not the bottleneck for the previous experiment, and secondly that not just the number of the encoding neurons but also their firing rate can be a limiting factor for the real-time performance.

Figure 3**b** shows the latency of the toolchain, i.e. the difference in time between a change in sensory input and

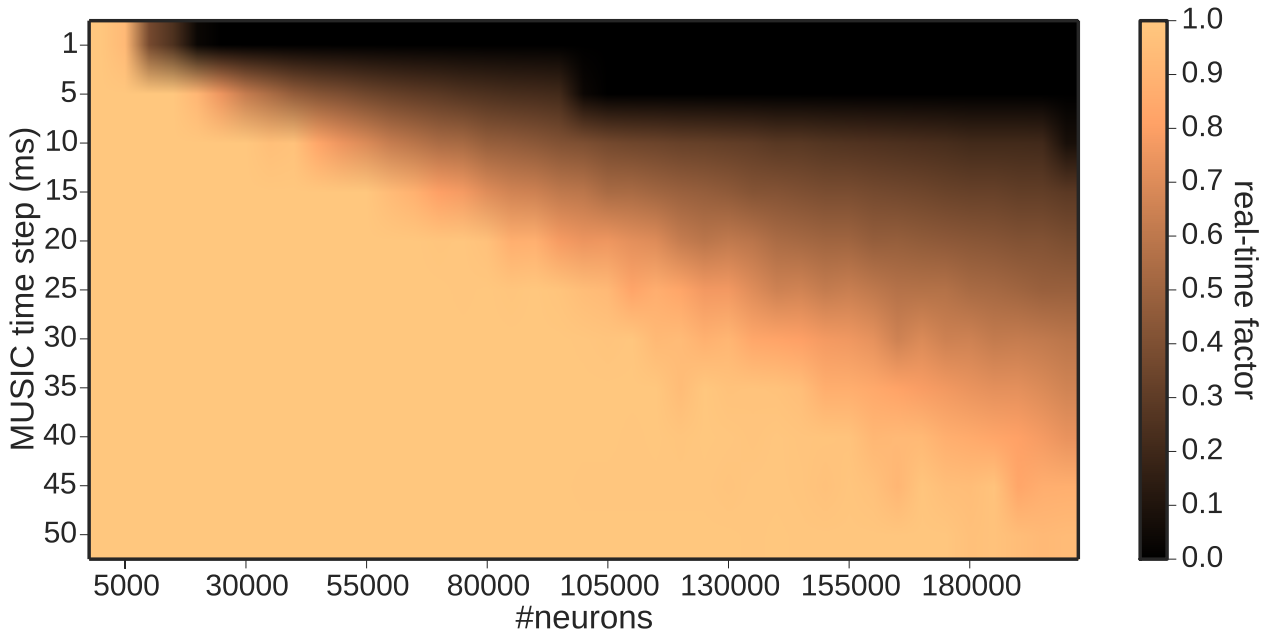

Figure 4: Real-time factor as a function of the MUSIC time step and number of neurons.

receiving a motor command evoked by that change (see Section 2.3.1). The latency increases linearly with the MUSIC time step; a MUSIC time step of $1\,\mathrm{ms}$, results in a latency of about $70\,\mathrm{ms}$, growing to around $350\,\mathrm{ms}$ for a time step of $50\,\mathrm{ms}$ when using a neural simulator. In general, the presence of a neural simulator in the toolchain increases the latency slightly, independently of which simulator is chosen.

Figure 4 depicts the dependency of real-time capability of the toolchain on the MUSIC time step. The border of real-time simulation capability increases linearly with the MUSIC time step, from $10,000$ neurons for a time step of $1\,\mathrm{ms}$ to $185,000$ neurons for a $50\,\mathrm{ms}$ time step. However, as demonstrated in Figure 3b, the latency also increases with the MUSIC time step. From these results, we conclude that the latency and the dimensionality of the input are two conflicting properties, which have to be balanced for the specific use case.

3.1 Implementation of a Braitenberg Vehicle

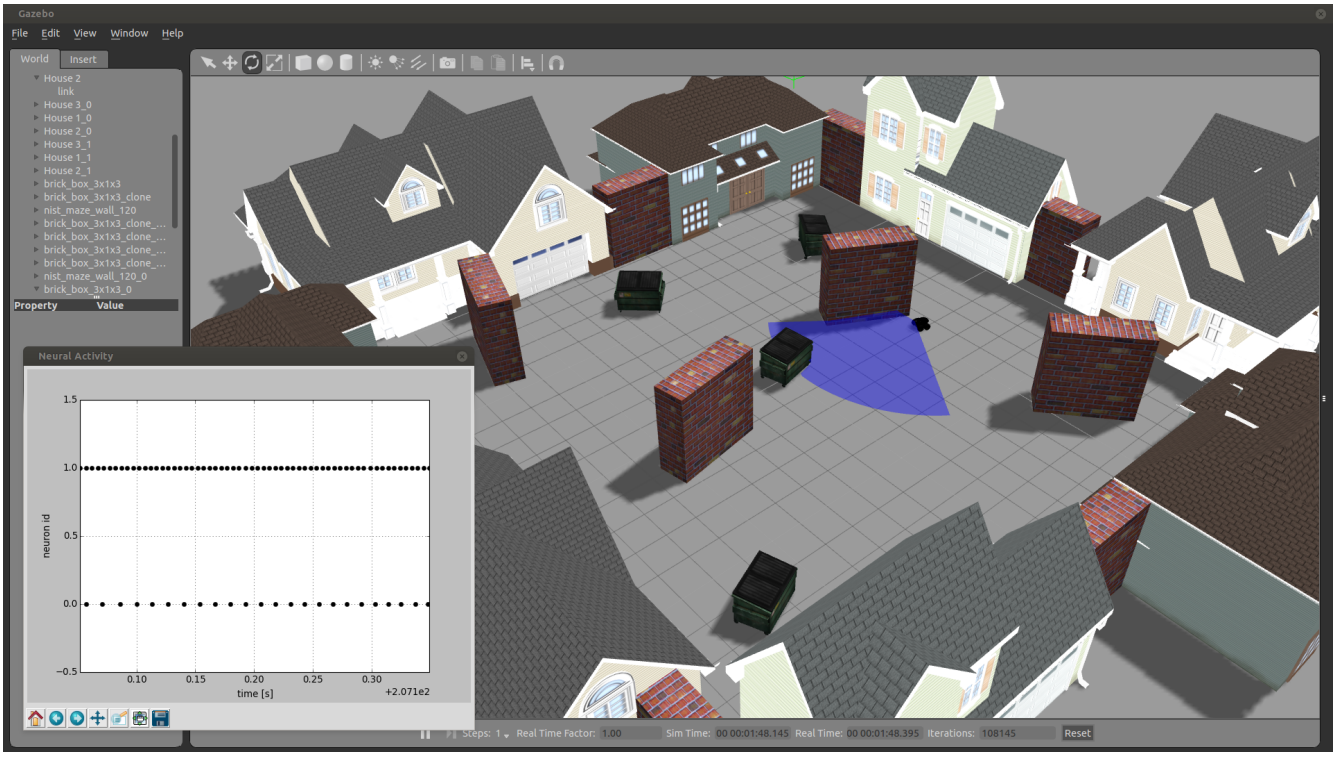
Gazebo
File Edit View Window Help
World Insert
Neural Activity
neuron id
time [s]
Steps: 1 Real Time Factor: 1.00 Sim Time: 00 00:01:48.145 Real Time: 00 00:01:48.395 Iterations: 108145 Reset

Figure 5: A Braitenberg Vehicle simulated in Gazebo, controlled by simulated neurons in NEST with the use of the ROS-MUSIC toolchain. The inset shows the spiking activity of the two controlling neurons.

As an example for the usage of the toolchain, we created a Braitenberg Vehicle III "Explorer" (Braitenberg, 1986) which is simulated in the robotic environment Gazebo (see Figure 5). The Braitenberg Vehicle is implemented as a four-wheeled mobile robot with an attached laser scanner for sensory input. With the use of the ROS-MUSIC toolchain, two neurons, simulated in NEST, control this vehicle to avoid obstacles. A video of the demonstration can be found in the supplementary material and the source code is available on GitHub[7] .

4 Discussion

We have described our plug-in for MUSIC, which allows any neuronal network simulator implementing a MUSIC interface to communicate with any robotics simulators implementing a ROS interface. The plug-in converts continuous signals

[7] https://github.com/weidel-p/ros_music_adapter

from a robotics simulator into spike train signals for a neuronal network simulator and vice versa. We showed that the toolchain allows real-time performance for a wide range of configurations and provided a simple working example. In the following, we discuss limits and perspectives for this approach.

Performance

Dependent on the choice of encoder, our toolchain allows the simulation of around $20,000$ neurons with a NEF encoder and up to $150,000$ neurons with a regular rate or Poisson encoder and achieve real-time performance. Using the NEF encoder, it is recommended to have at least 100 neurons per input dimension in order to encode and decode stimuli with root-mean-square errors of less than 1% (Eliasmith & Anderson, 2004), meaning that our toolchain can encode a 200-dimensional input with NEF in real time on a single process of our hardware.

If a higher dimensional input is required, there are two solutions at hand. First, the NEF encoder could be parallelized and run in more than one process. Second, if a less sophisticated but computationally cheaper way of encoding is acceptable for the scientific question at hand, we also provide regular rate and Poissonian encoding mechanisms. A previous neurorobotic interface implemented the communication between ROS and the neural simulators on the Python level (Hinkel et al., 2015). Although no performance data were published in that study, it is clear that the performance would be strongly limited by this approach.

Another critical characteristic of the toolchain is the latency between sensory input and motor command output of the toolchain, or in other words the reaction time. To minimize the overhead and avoid repeated communication of the same data, it makes sense to set the MUSIC time step equal to the sensory update or motor command rate. However, Figure 3 demonstrates that this leads to a rather high latency. The architecture of communicating processes currently carries an unavoidable source of latency, since data is buffered in MUSIC for every pair of communicating processes. In any chain of communicating processes, the latency due to buffering will thus grow in proportion to the length of the chain. In future work, this can be tackled by combining multiple adapter, encoder and decoder steps in a single process, thus simultaneously minimizing communication and reaching a lower latency, and/or by introducing non-buffered inter-process communication. Combining adapters, encoders and decoders can be realised using a plug-in adapter API which maintains independence between the software components.

Complexity of neuronal simulations

In our measurements of the toolchain incorporating NEST and NEURON, we used only minimal networks solving no particular task. However, neural network models can rapidly become computationally complex, especially when incorporating

synaptic plasticity, large network sizes, or multi-compartment neuron models. Such models would be impossible to run in real-time using the current approach. In some cases this issue could be solved by increasing the computational resources, using a cluster or supercomputer (Helias et al., 2012; Kunkel et al., 2014). However, this does not guarantee real-time execution, as neuronal network simulators do not have perfect scaling, due to serial components in their algorithms and communication overhead. Even if the problem can theoretically be addressed in real-time by increasing resources, it might not be feasible to access the quantity required. An alternative prospect is offered by neuromorphic hardware, i.e. hardware purpose built for simulation/emulation of neuronal networks. Two examples that are being developed within the framework of the Human Brain Project are SpiNNaker (Furber et al., 2013) and NM-PM1 (Schemmel et al., 2010), which have the potential to speed up the neuronal simulation massively. In particular SpiNNaker has great potential for neurorobotic applications and an interface between SpiNNaker and MUSIC is already in the prototype phase. This development enhances the value of the toolchain we describe for the neurorobotic community.

Whereas the real-time property is important for robotic control, it is not nearly so important for addressing questions in the field of computational neuroscience. Here, the advantage of the toolchain is that a neuronal network simulation can be provided with rich sensory input from an agent interacting with an environment that is easy for the researcher to configure. In this case, arbitrarily computationally demanding networks can be coupled with robotics simulators simply by slowing down the latter to compensate - gazebo, for example, provides a parameter to conveniently control the time scaling.

Applications

The toolchain we describe gives researchers in computational neuroscience the possibility to test their hypothesis and models under more realistic conditions of noisy sensory input, and researchers in neurorobotics the opportunity to investigate more realistic neurally based controllers. One area of potential interest is the ability to construct interactions between robotic simulators and hybrid neuronal simulations on multiple scales, e.g. a network of point neuron models into which detailed biophysical models are embedded, simulated by NEST and NEURON respectively. This demonstrates a further advantage of the MUSIC-based interaction over pairings of particular simulators or the Python-based interaction presented by Hinkel et al. (2015).

Moreover, our toolchain is particularly well suited for studying closed-loop scenarios, where the neural network receives stimuli from a complex environment and produces an output, which in turn causes the robotic agent to perform actions within that environment. For example, a robotic agent can be placed in a classic experimental set-up like a T-maze and the behaviour of the robot adapted by a neurally implemented reinforcement learner (Jitsev et al., 2012; Potjans et al., 2011; Frémaux et al., 2013; Friedrich & Lengyel, 2015). Here, there is a clear advantage over studying such

questions just using neural simulators, as the representation of an external environment as a collection of neural recorders and stimulators is complex, and difficult to either generalize or customize. By separating the concerns of environmental, motor, and sensory representation from those of neural processing, our toolchain provides a highly flexible and performant research approach.

Acknowledgments

We thank Karolína Korvasová, Tom Tetzlaff and Jenia Jitsev for helpful discussions, and Jenia also for his comments on an earlier version of this manuscript. We acknowledge partial support by the German Federal Ministry of Education through our German-Japanese Computational Neuroscience Project (BMBF Grant 01GQ1343), EuroSPIN, the Helmholtz Alliance through the Initiative and Networking Fund of the Helmholtz Association and the Helmholtz Portfolio theme "Supercomputing and Modeling for the Human Brain" and the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 604102 (HBP). We also thank the INCF for travel support enabling work meetings.

References

Antonelo, E. A., Schrauwen, B., Dutoit, X., Stroobandt, D., & Nuttin, M. (2007). Event Detection and Localization in Mobile Robot Navigation Using Reservoir Computing. In J. M. de Sá, L. A. Alexandre, W. Duch, & D. Mandic (Eds.), *Artificial Neural Networks – ICANN 2007*, Volume 4669 of *Lecture Notes in Computer Science*, Berlin, Heidelberg, pp. 660–669. Springer Berlin Heidelberg.

Averbeck, B. B. (2009). Poisson or not Poisson: differences in spike train statistics between parietal cortical areas. *Neuron 62*(3), 310–1.

Braitenberg, V. (1986). *Vehicles: Experiments in synthetic psychology*. MIT press.

Calvin, W. H., & Stevens, C. F. (1967). Synaptic noise as a source of variability in the interval between action potentials. *Science (New York, N.Y.) 155*(3764), 842–4.

Chorley, P., & Seth, A. K. (2008). Closing the sensory-motor loop on dopamine signalled reinforcement learning. In *From Animals to Animats 10*, pp. 280–290. Springer.

Dasgupta, S., Wörgötter, F., & Manoonpong, P. (2013). Information dynamics based self-adaptive reservoir for delay temporal memory tasks. *Evolving Systems 4*(4), 235–249.

Dayan, P., & Abbott, L. F. (2001). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*.

Djurfeldt, M., Hjorth, J., Eppler, J. M., Dudani, N., Helias, M., Potjans, T. C., Bhalla, U. S., Diesmann, M., Kotaleski, J. H., & Ekeberg, Ö. (2010). Run-time interoperability between neuronal network simulators based on the music framework. *Neuroinformatics 8*(1), 43–60.

Eliasmith, C., & Anderson, C. H. (2004). *Neural engineering: Computation, representation, and dynamics in neurobiological systems*. MIT press.

Eppler, J. M., Pauli, R., Peyser, A., Ippen, T., Morrison, A., Senk, J., Schenck, W., Bos, H., Helias, M., Schmidt, M., Kunkel, S., Jordan, J., Gewaltig, M.-O., Bachmann, C., Schuecker, J., Albada, S., Zito, T., Deger, M., Michler, F., Hagen, E., Setareh, H., Riquelme, L., Shirvani, A., Duarte, R., Deepu, R., & Plesser, H. E. (2015). Nest 2.8.0.

Frémaux, N., Sprekeler, H., & Gerstner, W. (2013). Reinforcement Learning Using a Continuous Time Actor-Critic Framework with Spiking Neurons. *PLoS Computational Biology 9*(4), e1003024.

Frémaux, N., Sprekeler, H., & Gerstner, W. (2013). Reinforcement learning using a continuous time actor-critic framework with spiking neurons. *PLoS Comput Biol 9*(4), e1003024.

Friedrich, J., & Lengyel, M. (2015). Goal-directed decision making with spiking neurons.

Furber, S. B., Lester, D. R., Plana, L. A., Garside, J. D., Painkras, E., Temple, S., & Brown, A. D. (2013). Overview of the spinnaker system architecture. *Computers, IEEE Transactions on 62*(12), 2454–2467.

Gewaltig, M.-O., & Diesmann, M. (2007). Nest (neural simulation tool). *Scholarpedia 2*(4), 1430.

Hanuschkin, A., Kunkel, S., Helias, M., Morrison, A., & Diesmann, M. (2010). A general and efficient method for incorporating precise spike times in globally time-driven simulations. *Frontiers in neuroinformatics 4*(October), 113.

Helias, M., Kunkel, S., Masumoto, G., Igarashi, J., Eppler, J. M., Ishii, S., Fukai, T., Morrison, A., & Diesmann, M. (2012). Supercomputers ready for use as discovery machines for neuroscience. *Front Neuroinform 6*(26), 2.

Hines, M., & Carnevale, N. T. (2001). Neuron: a tool for neuroscientists. *The Neuroscientist 7*(2), 123–135.

Hinkel, G., Groenda, H., Vannucci, L., Denninger, O., Cauli, N., & Ulbrich, S. (2015). A Domain-Specific Language (DSL) for Integrating Neuronal Networks in Robot Control. In *Proceedings of the 2015 Joint MORSE/VAO Workshop on Model-Driven Robot Software Engineering and View-based Software-Engineering*, MORSE/VAO '15, New York, NY, USA, pp. 9–15. ACM.

Jitsev, J., Morrison, A., & Tittgemeyer, M. (2012). Learning from positive and negative rewards in a spiking neural network model of basal ganglia. In *Neural Networks (IJCNN), The 2012 International Joint Conference on*, pp. 1–8. IEEE.

Kunkel, S., Schmidt, M., Eppler, J. M., Plesser, H. E., Masumoto, G., Igarashi, J., Ishii, S., Fukai, T., Morrison, A., Diesmann, M., et al. (2014). Spiking network simulation code for petascale computers. *Frontiers in neuroinformatics 8*(78).

Legenstein, R., & Maass, W. (2014). Ensembles of spiking neurons with noise support optimal probabilistic inference in a dynamically changing environment. *PLoS Comput Biol 10*(10), e1003859.

Moren, J., Sugimoto, N., & Doya, K. (2015). Real-time utilization of system-scale neuroscience models. *Journal of the Japanese Neural Network Society 22*(3), 125–132.

Morrison, A., Aertsen, A., & Diesmann, M. (2007). Spike-timing-dependent plasticity in balanced random networks. *Neural computation 19*(6), 1437–67.

Potjans, W., Diesmann, M., & Morrison, A. (2011). An imperfect dopaminergic error signal can drive temporal-difference learning. *PLoS Comput Biol 7*(5), e1001133.

Quiñonez, Y., Ramirez, M., Lizarraga, C., Tostado, I., & Bekios, J. (2015). Autonomous Robot Navigation Based on Pattern Recognition Techniques and Artificial Neural Networks. In J. M. F. Vicente, J. R. Álvarez-Sánchez, F. d. l. P. López, F. J. Toledo-Moreo, & H. Adeli (Eds.), *Bioinspired Computation in Artificial Systems*, Number 9108 in Lecture Notes in Computer Science, pp. 320–329. Springer International Publishing.

Quigley, M., Conley, K., Gerkey, B., Faust, J., Foote, T., Leibs, J., Wheeler, R., & Ng, A. Y. (2009). Ros: an open-source robot operating system. In *ICRA workshop on open source software*, Volume 3, pp. 5.

Rotter, S., & Diesmann, M. (1999). Exact digital simulation of time-invariant linear systems with applications to neuronal modeling. *Biological cybernetics 81*(5-6), 381–402.

Schemmel, J., Bruderle, D., Grubl, A., Hock, M., Meier, K., & Millner, S. (2010). A wafer-scale neuromorphic hardware system for large-scale neural modeling. In *Circuits and systems (ISCAS), proceedings of 2010 IEEE international symposium on*, pp. 1947–1950. IEEE.